Large-scale JPEG image steganalysis using hybrid deep-learning framework

Jishen Zeng, *Student Member, IEEE,* Shunquan Tan*, *Member, IEEE,* Bin Li, *Member, IEEE,* and Jiwu Huang, *Fellow, IEEE*

arXiv:1611.03233v2 [cs.MM] 26 Jan 2017

***Abstract*—Deep learning frameworks have recently achieved superior performance in many pattern recognition problems. However, adoption of deep learning in image steganalysis is still in its initial stage. In this paper we propose a hybrid deep-learning framework for JPEG steganalysis incorporating the domain knowledge behind rich steganalytic models. We prove that the convolution phase and the quantization & truncation phase of the rich models are not learnable in deep convolutional neural networks. Based on theoretical analysis, our proposed framework involves two main stages. The first stage is hand-crafted, corresponding to the convolution phase and the quantization & truncation phase of the rich models. The second stage is a compound deep neural network containing three deep subnets in which the model parameters are learned in the training procedure. By doing so, we ably combine some merits of rich models into our proposed deep-learning framework. We have conducted extensive experiments on a large-scale dataset extracted from ImageNet. The primary dataset used in our experiments contains 500,000 cover images, while our largest dataset contains five million cover images. Our experiments show that the proposed framework outperforms all other state-of-the-art steganalytic models either hand-crafted or learned using deep networks in the literature. Furthermore, we demonstrate that our framework is insensitive to JPEG blocking artifact alterations and the learned model can be easily transferred to a different attacking target. Both of these properties are of critical importance in practical applications. According to our best knowledge, This is the first report of deep learning in image steganalysis validated with large-scale test data.**

***Index Terms*—hybrid deep-learning framework, CNN network, steganalysis, steganography.**

I. Introduction

STEGANOGRAPHY is a technique of covert communications. As an adversary of steganography, steganalysis tries to detect the presence of secret messages [1]. Digital images are among the most best cover media due to their ease of acquisition and wide distribution in the Internet. As a result, image steganography and steganalysis has become an active topic in recent years.

Image steganography can be divided into two main categories: spatial-domain and frequency-domain steganography. The latter focuses primarily on JPEG images due to their ubiquitous nature. Both categories in state-of-the-art algorithms adopt content-adaptive embedding schemes [2], [3]. Most of these schemes use an additive distortion function defined as the sum of embedding costs of all changed elements. From early HUGO [4] and WOW [5], to latest HILL [6], [7], MVG [8], [9] and MiPOD [10], the past few years witnessed the flourish of additive schemes in spatial domain.[1] In JPEG domain, UED [11] and UERD [12] are two additive schemes with good security performance. UNIWARD proposed in [13] is an additive distortion function which can be applied for embedding both in spatial and JPEG domains. Its JPEG version, J-UNIWARD, achieves best performance [12], [13]. Research on non-additive distortion functions has made great progress in the spatial domain [14], [15]. However, analogous schemes have not yet been proposed in the JPEG domain. Although utilizing side information of a pre-cover image (raw or uncompressed) can improve the security of JPEG steganography [12], [13], its applicability remains limited due to scarce availability of pre-cover images.

Steganalysis, on the other hand, has evolved from early targeted steganalytic algorithms into advanced universal steganalytic frameworks. Most of modern universal steganalytic detectors use a rich model with tens of thousands of features [16]–[19] and an ensemble classifier [20]. In spatial domain, SRM [16] and its selection-channel-aware variants [17]–[19] reign supreme. In JPEG domain, DCTR [21] feature set combines relatively low dimensionality and competitive performance, while PHARM [22] exhibits better performance, although at the cost of higher dimensionality w.r.t. DCTR. SCA proposed in [23] is a selection-channel-aware variant of JPEG rich models targeted at content-adaptive JPEG steganography.

In recent years, with help of parallel computing accelerated by GPU (Graphics Processing Unit) and huge amounts of training data, deep learning frameworks have achieved overwhelming superiority over conventional approaches in many pattern recognition and machine learning problems [24]. Researchers in image steganalysis have also tried to investigate the potential of deep learning frameworks in this field [25]–[27]. Tan et al. explored the application of stacked convolutional auto-encoders, a specific form of deep learning frameworks in image steganalysis [25]. Qian et al. proposed a steganalyzer based on CNN (Convolutional Neural Network), achieving performance close to SRM [26]. Xu et al. constructed another CNN-based steganalyzer [27] equipped with

This work was supported in part by the NSFC (61332012, 61402295, 61572329), Guangdong NSF (2014A030313557). *(Corresponding author: Shunquan Tan.)*

S. Tan is with College of Computer Science and Software Engineering, Shenzhen University. J. Zeng, B. Li, and J. Huang are with College of Information Engineering, Shenzhen University.

All the members are with Shenzhen Key Laboratory of Media Security, Guangdong Province, 518060 China. (e-mail: tansq@szu.edu.cn).

[1]Throughout this paper, the acronyms used for the steganographic and steganalytic algorithms are taken from the original papers. The corresponding full names are omitted for brevity.

BN (Batch Normalization) layers [28]. Its performance slightly surpass SRM. However, all of the above approaches [25]–[27], focusing on spatial-domain steganalysis, are all evaluated on the BOSSBase (v1.01) dataset [29]. BOSSBase is arguably not representative of real-world steganalysis performance [30]. With only 10,000 images, deep learning frameworks trained on BOSSBase are prone to overfitting. Furthermore, none of existing studies addressed the application of deep learning frameworks in JPEG steganalysis.

In this paper, we proposed a hybrid deep-learning framework for large-scale JPEG steganalysis. We firstly prove that the two crucial components of rich models, i.e. the convolutional kernels used to extract diverse noise residuals and the threshold quantizers used to reduce model complexity, cannot be efficiently learned by CNNs. Based on theoretical analysis, our proposed framework combines the bottom hand-crafted convolutional kernels and threshold quantizers pairs with the upper compact deep-learning model. Extensive experiments conducted on a large-scale dataset extracted from ImageNet [31] show that the performance of our proposed framework is superior to both rich-model-based steganalyzers and other deep-learning-based steganalyzers in the literature. Furthermore, the proposed hybrid deep-learning framework exhibits good properties, including good transfer ability among different steganographic algorithms and insensitivity to the alteration of JPEG blocking artifact.

The rest of the paper is organized as follows. Sect. II provides a theoretical foundation of our approach and then describes the proposed hybrid deep-learning framework in detail. Results of experiments conducted on ImageNet images are presented in Sect. III. Finally, we make a conclusion in Sect. IV.

II. Our proposed JPEG steganalytic framework

In this section, we firstly provide an overview of rich models for JPEG steganalysis and the training procedure of CNNs as preliminaries. Then we prove two propositions, which act as the theoretical foundation of this paper. Finally we present the conceptual architecture or our proposed hybrid deep-learning framework with two types of subnet configurations.

A. Preliminaries

State-of-the-art rich-models for JPEG steganalysis [21]–[23] take decompressed JPEG images as input. Like in their spatial-domain counterparts, the feature extraction procedure of JPEG rich models can be divided into three phases:

- *Convolution*: The target image is convolved with a set of kernels to generate diverse noise residuals. The purpose of this phase is to suppress the image contents as far as possible.
- *Quantization and truncation* (Q&T): Different quantized and truncated versions of each residual are calculated to further improve diversity of resulting features, as well as reduce the computational complexity.
- *Pooling*: Aggregation of values in noise residuals to further reduce feature dimensionality.

Take DCTR [21] for example. Given a $M \times N$ JPEG image, it is firstly decompressed to the corresponding spatial-domain version $\mathbf{X} \in \mathbb{R}^{M \times N}$. Sixty-four 8×8 DCT basis patterns are defined as $\mathbf{B}^{(k,l)} = (B_{mn}^{(k,l)}), 0 \le k, l \le 7, 0 \le m, n \le 7$:

$$B_{mn}^{(k,l)} = \frac{w_k w_l}{4} cos\frac{\pi k(2m+1)}{16} cos\frac{\pi l(2n+1)}{16}, \tag{1}$$

where $w_0 = \frac{1}{\sqrt{2}}$, $w_k = 1$ for $k > 0$. $\mathbf{X}$ is convolved with $\mathbf{B}^{(k,l)}$ to generate 64 noise residuals $\mathbf{U}^{(k,l)}, 0 \le k, l \le 7$:

$$\mathbf{U}^{(k,l)} = \mathbf{X} * \mathbf{B}^{(k,l)}, \tag{2}$$

Then the elements in each $\mathbf{U}^{(k,l)}$ are quantized with quantization step q and truncated to a threshold T. The DCTR features are constructed based on certain pooling operation that collect specific first-order statistics of the absolute values of the quantized and truncated elements in each $\mathbf{U}^{(k,l)}$.

In [25], we pointed out that the above structure of rich models resembles CNNs, which are a cascade of alternating convolutional layers, regulation layers (e.g. BN layers [28]) and pooling layers. Regardless of the types of the layers, they are made of units with learnable weights and biases. Each unit receives inputs from some units of a previous layer, performs a dot product with weights and optionally follows it with a nonlinear point-wise activation function. CNNs can be trained using backpropagation. The cascade of layers in a given CNN is denoted as $[L_1, L_2, \cdots, L_n]$, where L_1 is the input layer and L_n is the output layer. Let $a_i^{(l)}$ denote the activation (output) of unit i in layer L_l. For L_1, $a_i^{(1)}$ is the i-th input. $W_{ij}^{(l)}$ denotes the weight associated with unit i in L_l and unit j in L_{l+1}, while $b_j^{(l)}$ denotes the bias associated with unit j in L_{l+1}. The weighted sum of inputs to unit j in L_{l+1} is defined as:

$$z_j^{(l+1)} = \sum_i W_{ij}^{(l)} a_i^{(l)} + b_j^{(l)} \tag{3}$$

and $a_j^{(l+1)} = f(z_j^{(l+1)})$ where $f(\cdot)$ is the activation function. The set of all $W_{ij}^{(l)}$ and $b_j^{(l)}$ constitutes the parameterization of a neural network and is denoted as W and b, respectively. W and b constitute the parameter set of a neural network. For a mini-batch of training features-label pairs $\{(x^{(1)}, y^{(1)}), \cdots, (x^{(m)}, y^{(m)})\}$, the goal of backpropagation is to minimize the overall cost function $J(W, b)$ with respect to W and b:

$$J(W, b) = \frac{1}{m} \sum_{h=1}^{m} J(W, b; x^{(h)}, y^{(h)}) + R(W) \tag{4}$$

where $R(W)$ is a regularization term and $J(W, b; x^{(h)}, y^{(h)})$ is an error metric with respect to a single example $(x^{(h)}, y^{(h)})$. For each training sample, the backpropagation algorithm firstly performs a feedforward pass and computes the activations for layers L_2, L_3 and so on up to the output layer L_n. Then it propagates partial derivatives from the output layer L_n back to the second last layer L_2. The partial derivatives with respect to $W_{ij}^{(l)}$ and $b_j^{(l)}$, $l = n, n-1, \cdots, 2$ are calculated as:

$$\begin{cases} \frac{\partial}{\partial W_{ij}^{(l)}} J(W, b; x^{(h)}, y^{(h)}) = a_i^{(l)} \vartheta_j^{(l+1)} \ , \\ \frac{\partial}{\partial b_j^{(l)}} J(W, b; x^{(h)}, y^{(h)}) = \vartheta_j^{(l+1)} \ , \end{cases} \tag{5}$$

where

$$\begin{cases} \vartheta_j^{(n)} = \frac{\partial}{\partial a_j^{(n)}} J(W,b;x^{(h)},y^{(h)}) f'(z_j^{(n)}) \ , \\ \vartheta_j^{(l)} = (\sum_k W_{jk}^{(l)} \vartheta_k^{(l+1)}) f'(z_j^{(l)}) \ . \end{cases} \tag{6}$$

Gradient descent is used to find the optimal W and b. In the optimization procedure, it updates W and b according to steps proportional to the negative of the average of m gradients each of which is the vector whose components are the partial derivatives in (5) [32].

B. Theoretical foundations of our proposed framework

Since rich steganalytic models exhibit structure similarity to CNNs [25], [27], one interesting question is whether it is possible for CNNs to utilize the domain knowledge behind rich models, especially the specific kernel matrices in the convolutional phase and the Q&T phase. In the following propositions, we provide a negative argument and claim that not only the convolution phase but also the Q&T phase of rich models are not learnable for CNNs.

Proposition 1.

1) *Kernel matrices used to generate noise residuals in rich models cannot be evolved from randomly initialized kernels of a convolutional layer in CNN. This problem is unsolvable even with traditional data preprocessing strategies.*
2) *Better performing kernels cannot be evolved from the kernels possessing the same parameters as those used in rich models, provided the model is trained with gradient descent.*

Proof: State-of-the-art steganalytic feature extractors, either in spatial domain or in JPEG domain, take the spatial representation of target image as input [16]–[19], [21]–[23]. We follow this approach in our research. Therefore, a grayscale input image can be represented as $\mathbf{X} = (x_{pq})^{M\times N} = \mathbf{C} + \mathbf{N}$, where $\mathbf{C} = (c_{pq})^{M\times N}$, $c_{pq} \in [0, 255]$ denotes the corresponding cover image, and $\mathbf{N} = (n_{pq})^{M\times N}$ denotes the additive stego noise. When $\mathbf{X}$ is an innocent cover image, $\mathbf{N}$ is a zero matrix. When $\mathbf{X}$ is a spatial-domain stego image generated by a ternary embedding scheme, $n_{pq} \in \{-1, 0, 1\}$ denotes the zero mean stego noise. When $\mathbf{X}$ is a JPEG-domain stego image, the situation is comparatively complicated due to the fact that the zero mean ±1 stego noise is directly added to quantized DCT coefficients.

Consider a JPEG target image $\widetilde{\mathbf{X}}$ represented as $\widetilde{\mathbf{X}} = \widetilde{\mathbf{C}} + \widetilde{\mathbf{N}}$, where $\widetilde{\mathbf{C}}$ denotes the corresponding quantized DCT coefficients and $\widetilde{\mathbf{N}} = (\widetilde{n}_{pq})^{M\times N}$ denotes the additive ±1 noises added to the DCT coefficients. Due to linearity property of the DCT/IDCT transform, the following equalities between $\{\mathbf{X},\ \mathbf{C},\ \mathbf{N}\}$ and $\{\widetilde{\mathbf{X}},\ \widetilde{\mathbf{C}},\ \widetilde{\mathbf{N}}\}$ hold: [2]

$$\mathbf{X} = \mathrm{IDCT}(\mathbf{Q} \circ \widetilde{\mathbf{X}}),\ \mathbf{C} = \mathrm{IDCT}(\mathbf{Q} \circ \widetilde{\mathbf{C}}),\ \mathbf{N} = \mathrm{IDCT}(\mathbf{Q} \circ \widetilde{\mathbf{N}}) \tag{7}$$

where $\mathbf{Q} = (q_{pq})^{M\times N}$ is the quantization matrix and $\circ$ represents the element-wise product. That is to say, the corresponding stego noise in the spatial-domain representation of the decompressed JPEG image is still additive and is a linear mapping of $\widetilde{\mathbf{N}}$. Therefore due to the energy conservation principle:

$$\sum_{p=1}^{M}\sum_{q=1}^{N} |n_{pq}|^2 = \sum_{p=1}^{M}\sum_{q=1}^{N} |q_{pq} \cdot \widetilde{n}_{pq}|^2 \tag{8}$$

in which Frobenius norm is used to measure the energy of the noise.

In order not to introduce visible and statistically detectable artifacts, state-of-the-art JPEG steganographic algorithms (e.g. J-UNIWARD and UERD) try to modify only low-frequency non-zero coefficients. As a result, firstly the proportion of non-zero ± 1 $\widetilde{n}_{pq}$ with respect to entire DCT coefficients is small, and secondly the magnitudes of q_{pq} multiplies to each non-zero $\widetilde{n}_{pq}$ is relatively small since low-frequency coefficients are assigned small quantization steps. Therefore the total energy of the right-hand side of (8) is kept as small as possible by state-of-the-art JPEG steganographic algorithms. Furthermore, since the elements of $\widetilde{\mathbf{N}}$ are uncorrelated, the IDCT transform spreads the energy of the noise among the elements of $(n_{pq})^{M\times N}$ in the left-hand side of (8). As a consequence, the magnitude of most of the elements of $\mathbf{N}$ matrix remain tiny with respect to the corresponding elements of $\mathbf{C}$ even for a stego JPEG image with high embedding rate.

In order to verify the above conclusion experimentally, we generated 50,000 J-UNIWARD stego images with 0.4 bpnzAC (bits per non-zero cover AC DCT coefficient) from basic50K (see Sect. III-A) dataset. We only considered their spatial representations. Our major concern was the ratio of $|c_{pq}|$, the magnitude of a cover pixel in spatial domain, to $|n_{pq}|$, the corresponding stego noise in spatial domain introduced by J-UNIWARD. The average of the frequency distributions of ratios of $|c_{pq}|$ to rounded $|n_{pq}|$ for those n_{pq} with non-zero rounded absolute values in the 50,000 stego images is shown in Fig. 1. From Fig. 1 we can see that on average the magnitudes of the cover pixels are much larger than the corresponding stego noises introduced by J-UNIWARD in spatial domain. The mean of the average frequency distribution is 88.04, which implies that $|c_{pq}|$ is on average close to two orders of magnitude larger than $|n_{pq}|$ even with an embedding rate as high as 0.4 bpnzAC.

Suppose that we want to train a convolutional layer with kernels of size of $m\times n$ to generate noise residuals, it should be located at the second layer of the CNN hierarchy. Convolution is just a dot product with local-connected-and-shared weights. By rewriting (3) using two-dimensional indexing, setting $l = 1$ and restrict the size of the dot product to $m \times n$, we get:

$$\begin{aligned} z_j^{(2)} &= \sum_{p=1}^{M}\sum_{q=1}^{N} W_{pq,j}^{(1)} a_{pq}^{(1)} + b_j^{(1)} = \sum_{p=1}^{m}\sum_{q=1}^{n} W_{pq,j}^{(1)} a_{pq}^{(1)} + b_j^{(1)} \\ &= \sum_{p=1}^{m}\sum_{q=1}^{n} W_{pq,j}^{(1)} x_{pq} + b_j^{(1)} = \sum_{p=1}^{m}\sum_{q=1}^{n} W_{pq,j}^{(1)} (c_{pq} + n_{pq}) + b_j^{(1)} \\ &= \sum_{p=1}^{m}\sum_{q=1}^{n} W_{pq,j}^{(1)} c_{pq} + \sum_{p=1}^{m}\sum_{q=1}^{n} W_{pq,j}^{(1)} n_{pq} + b_j^{(1)} \end{aligned}$$

[2]The 8×8 DCT block structure in JPEG compression is neglected to simplify the discussion.

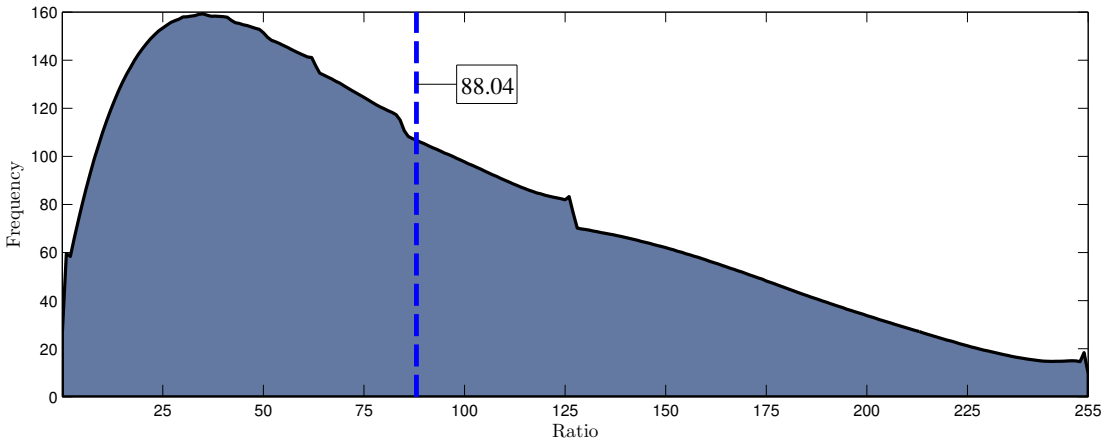

Fig. 1. Average frequency distribution of ratios of $|c_{pq}|$ to rounded $|n_{pq}|$ for those n_{pq} with non-zero rounded absolute values in 50,000 basic50K (see Sect. III-A) J-UNIWARD stego images with 0.4 bpnzAC. The vertical dashed line denotes the mean of the average frequency distribution.

$$\approx \sum_{p=1}^{m}\sum_{q=1}^{n} W^{(1)}_{pq,j} c_{pq} + b^{(1)}_j \tag{9}$$

In (9), $a^{(1)}_{kl} = x_{kl}$ since the convolutional layer is fed with input images. The approximation holds due to the discussed relation between the magnitudes of **C** and **N**. According to (5) and (6) we can see that in one iteration of backpropagation, when the gradient is backpropagated to L_2, $\frac{\partial}{\partial W^{(2)}_{jk}} J(W,b;x^{(h)},y^{(h)}) \propto a^{(2)}_j$ since $\vartheta^{(3)}_k$ is fixed. Then $a^{(2)}_j \propto z^{(2)}_j$ due to the fact that all of the commonly used activation functions, including Sigmoid, TanH and ReLU are monotonic increasing. As a result $\frac{\partial}{\partial W^{(2)}_{jk}} J(W,b;x^{(h)},y^{(h)}) \propto z^{(2)}_j$. According to (9), the influence of stego noise **N** to $\frac{\partial}{\partial W^{(2)}_{jk}} J(W,b;x^{(h)},y^{(h)})$ can be neglected. Therefore gradient descent algorithm will be guided by the cover image content rather than the stego noise. As a result, the kernel matrices used to generate noise residuals in rich models cannot be evolved from the CNN convolutional kernels.

Furthermore, since during the training procedure stego noise is mixed with image contents, traditional data preprocessing strategies, including mean subtraction, normalization, and whitening cannot narrow the gap of the order of magnitude between $|c_{pq}|$ and $|n_{pq}|$. Thus data preprocessing strategies cannot help resolving the problem. The proof of Part 1) is obtained.

Since the proof is independent of the kernel initialization of the corresponding convolutional layers, the influence of $\{n_{pq}\}$ to $\frac{\partial}{\partial W^{(2)}_{jk}} J(W,b;x^{(h)},y^{(h)})$ still can be neglected even when the convolutional kernels are initialized from the fillers in rich models, which concludes the proof of Part 2). ■

Please note that recently in a similar field, image forensics, Bayer et al. proposed a convolutional-layer regularizer which was claimed can be used to suppress the content of an image [33]. However, the regularization only affects the output residual maps of a given convolutional layer, rather than affects the magnitudes of the input image pixels. It cannot change the relation between the magnitudes of **C** and **N**. Therefore according to Part 2) of Proposition 1 we can see that the regularizer [33] just gives a hand-crafted constraint to the shape of the convolutional kernels from the very beginning. Training with the regularizer cannot result in better performing kernels, especially in the field of steganalysis.

Proposition 2.

1) *The Q&T non-linearity saturates and nullifies gradients and therefore cannot be put in the pipeline of gradient descent.*
2) *The Q&T in rich models can be neither evolved from commonly used activation functions in CNN, nor imitated by a cascade of layers in CNN.*

Proof: The Q&T in rich models takes noise residuals generated by kernel matrices as input, which can be modeled as:

$$a^{(2)}_j = f(z^{(2)}_i) = \begin{cases} \min([z^{(2)}_j/q], T) & \text{if } z^{(2)}_j >= 0 \\ \max([z^{(2)}_j/q], -T) & \text{if } z^{(2)}_j < 0 \end{cases} \tag{10}$$

where $z^{(2)}_j$ is an element of a given noise residual, $a^{(2)}_j$ is the corresponding activation output, q is the quantization step, $[\cdot]$ denotes the rounding operation, and T is a predefined threshold. (10) can be rewritten using a staircase function:

$$f(z^{(2)}_i) = \sum_{k=-T}^{T} k \cdot \chi_{A_k}(z^{(2)}_j/q) \tag{11}$$

where $k \in \mathbb{Z}$, χ_{A_k} is an indicator function of A_k and $\{A_k\}, k = -T, -T+1, \cdots, T-1, T$ is a union of intervals:

$$A_k = \begin{cases} (-\infty, T+0.5], & k = -T \\ (k-0.5, k+0.5], & -T < k < 0 \\ (-0.5, 0.5), & k = 0 \\ [k-0.5, k+0.5), & 0 < k < T \\ [T-0.5, +\infty), & k = T \end{cases} \tag{12}$$

Therefore the derivative is:

$$f'(z^{(2)}_j) = \frac{da^{(2)}_j}{dz^{(2)}_j} = \sum_{k=-T}^{T} k \cdot \delta\left(\frac{z^{(2)}_j}{q} - k + 0.5\right) \tag{13}$$

where $\delta(\cdot)$ is the Dirac delta function which is zero everywhere except at zero where it is infinite. As a result $f'(z^{(2)}_j)$ is zero along the entire domain of $z^{(2)}_j$ except the set of points $\{-T+0.5, -T+1.5, \cdots, T-1.5, T-0.5\}$. According to (6)

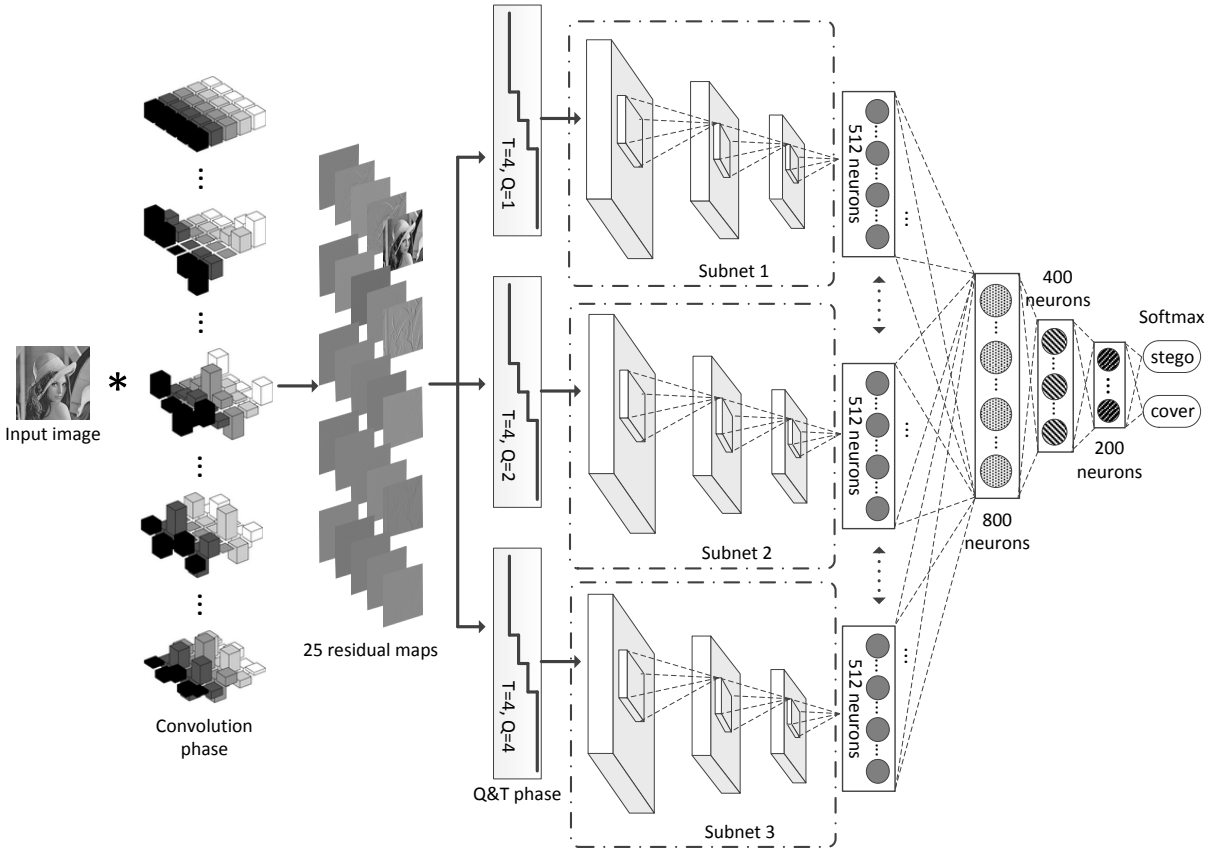

Fig. 2. Conceptual architecture of our proposed hybrid deep-learning framework. An input image is convolved with 5×5 DCT basis patterns to generate twenty-five corresponding residual maps. The residual maps are separately quantized and truncated with three different threshold quantizers. The central part of our framework includes three parallel CNN subnets. Each of them takes one quantized and truncated residual map as input, respectively. The features generated by them are concatenated and fed to a three-layer fully-connected neural network to yield the final prediction.

we can see that if (10) is used as the activation function, the partial derivative it passes on will vanish. That is to say the derivative does not exist if $z_j^{(2)}$ is located at one of the points in the set $\{-T+0.5, -T+1.5, \cdots, T-1.5, T-0.5\}$, otherwise the derivative is equal to zero. The corresponding gradient saturates if the partial derivative it passes on approaches to zero, and is nullified if there is no derivative. Thus the proof of Part 1) is obtained. Likewise, when a cascade of layers in CNN try to learn a staircase function represented in (11), the backpropagation will get stuck since the corresponding partial derivatives transmitted in the layers will approach to zero or infinite. Thus the proof of Part 2) is also obtained. ■

C. Details of our proposed hybrid deep-learning framework

Based on the theoretical analysis in Sect. II-B, we propose a hybrid deep-learning framework whose conceptual architecture is illustrated in Fig. 2. The proposed framework is composed of two stages. The first stage corresponds to the convolution phase and the Q&T phase of rich models. More specifically, this stage of our framework incorporated the first two phases of DCTR [21]. All model parameters in this stage are hand-crafted and gradient-descent-based learning is disabled. What makes this stage different from DCTR is that DCTR uses sixty-four 8×8 DCT basis patterns and only one Q&T combination, while our proposed approach contains twenty-five 5×5 DCT basis patterns which are defined as $\mathbf{B}^{(k,l)} = (B_{mn}^{(k,l)}), 0 \le k, l \le 5, 0 \le m, n \le 5$:

$$B_{mn}^{(k,l)} = \frac{w_k w_l}{5} cos\frac{\pi k(2m+1)}{10} cos\frac{\pi l(2n+1)}{10}, \quad w_0 = 1,\ w_k = \sqrt{2} \text{ for } k > 0. \quad (14)$$

and three Q&T combinations, namely $(T = 4, Q = 1)$, $(T = 4, Q = 2)$ and $(T = 4, Q = 4)$. Given an input image, the convolution phase outputs twenty-five residual maps. The residual maps pass through the Q&T phase. Three different groups of quantized and truncated residual maps are generated. They constitute the input of the second stage. The motivation behind the design of the first stage is that we agree with the concepts in rich models [16]: model diversity is crucial to the performance of steganalytic detectors. The model diversity of our proposed framework is represented in twenty-five DCT basis patterns in the hand-crafted convolutional layer and the three Q&T combinations that followed. There are total $25 \times 3 = 75$ sub-models in our proposed framework.

The second stage is a compound deep CNN network in which the model parameters are learned in the training procedure. The bottom of the second stage is composed of three independent subnets with identical structure. Each subnet corresponds to one group of quantized and truncated residual maps. They take the residual maps as input and generate three

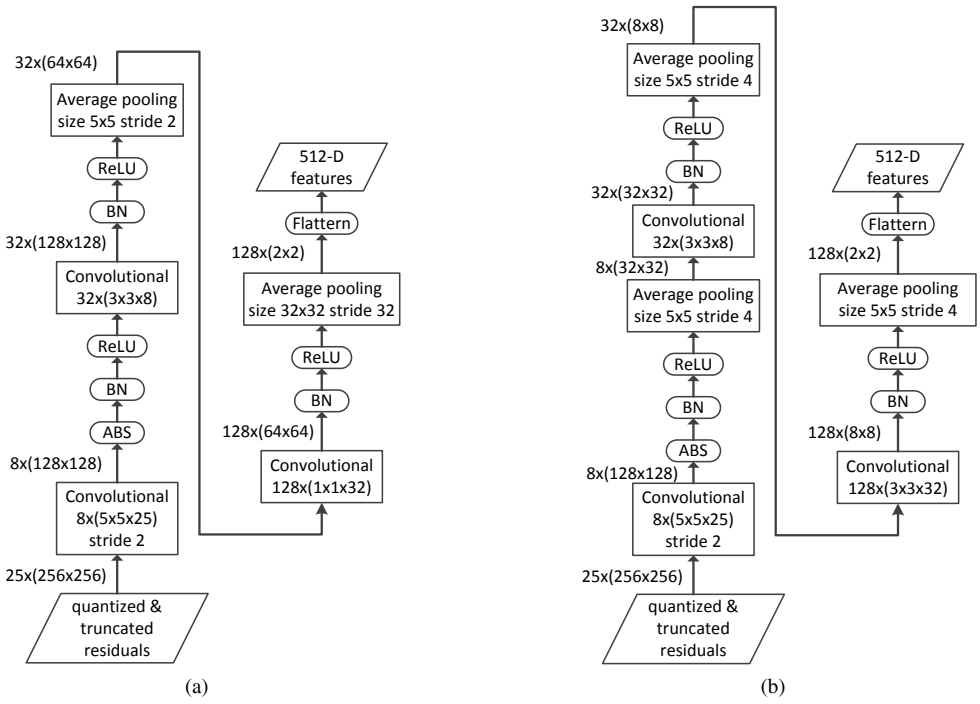

Fig. 3. Two types of subnet configurations. (a) *Type1* subnet. (b) *Type2* subnet.

feature vectors. As shown in Fig. 3, two types of subnet configurations are adopted in our proposed framework. Both of them contain three convolutional layers and output a 512-D (512 dimensional) feature vector. *Type1* subnet (Fig. 3(a)) adopts 1×1 convolutional kernels in the top-most convolutional layer and uses a single average pooling layer with large 32×32 pooling windows at the end, as suggested in Xu's model [27]. *Type2* subnet (Fig. 3(b)) is a traditional CNN configuration. Compared with the first configuration, it adopts progressive pooling layers and uses 3×3 convolutional kernels in the top-most convolutional layer. Due to the progressing pooling layers, *Type2* subnet is a relative GPU memory-efficient model. The GPU memory requirement of *Type2* subnet is only one-seventh of that of *Type1* subnet. Both configurations have in common are the BN layers which follow every convolutional layers.

The three 512-D feature vectors output by the bottom subnets are concatenated together to generate a single 1536-D feature vector. The feature vector is subsequently fed into a four-layer fully-connected neural network which makes the final prediction. The successive layers of the fully-connected network contain 800, 400, 200, and 2 neurons, respectively. ReLU activation functions are used in all three hidden layers. The final layer contains two neurons which denote "stego" prediction and "cover" prediction, respectively. Softmax function is used to output predicted probabilities.

Recent researches on deep-learning revealed that ensemble prediction with independently trained deep-learning models can significantly improve the performance [34]. We also use model ensemble in the final prediction. Five versions of our proposed deep-learning models are independently trained and their decision are combined with majority voting.

III. Experimental results

A. Experiment setups

We adopted ImageNet [31], a large-scale image dataset containing more than fourteen million JPEG images, to evaluate the steganalytic performance of our proposed hybrid deep-learning framework. All of the experiments were conducted on a GPU cluster with eight NVIDIA® Tesla® K80 dual-GPU cards. By considering the computation capacity, we restricted the size of the target images to 256×256. We randomly selected 50 thousand, 500 thousand and 5,000 thousand (namely 5 million) JPEG images with size larger than 256×256 from ImageNet. Their left-top 256×256 regions were cropped, converted to grayscale and then re-compressed as JPEG with quality factor 75. [3] The resulting images constituted the following three basic cover image datasets:

- basic50K: The small-scale dataset used in our experiments. By comparing the detection performance of our proposed framework on basic50K and basic500K (see below), we can highlight the superiority of our proposed framework in large-scale dataset.
- basic500K: The major dataset for most all of our experiments, including the verification experiments to determine hyper-parameters of our proposed framework.

[3]The original quality factors of ImageNet images are diverse. Out of 10 million ImageNet images with size larger than 256×256, there are more than 1.5 million images whose quality factors cannot be detected by ImageMagick utility "identify", and roughly 8.3 million images with diverse quality factors which are larger than 75. We uniformly converted the quality factors of the selected images to 75 due to the following two reasons: Firstly, all the reported experiments of previous works, including DCTR, PHARM, and SCA, are conducted on images with quality factor 75 and 95. And secondly, if the target quality factor is set to 95, then for a majority of the selected images, we need to elevate their quality factors which may introduce exploitable artifacts.

- basic5000K: The largest-scale dataset used in our experiments. Due to the limitation of computation capacity, we only conducted the experiments on stego images with 0.4 bpnzAC.

Our implementation was based on the publicly available Caffe toolbox [35] with our implemented hand-crafted convolutional layer (with 5×5 DCT basis patterns) and Q&T layer according to (10). Our proposed models were trained using mini-batch stochastic gradient descent with "step" learning rate starting from 0.01 (stepsize: 5000; weight_decay: 0.0005; gamma: 0.9) and a momentum fixed to 0.9. The batch size in the training procedure was 64 and the maximum number of iterations was set to 20×10^4.

J-UNIWARD [13], UERD [12] and UED [11], the three state-of-the-art JPEG steganographic schemes, were our attacking targets in the experiments. The default parameters of the three steganographic schemes were adopted in our experiments. 50% cover images were randomly selected from basic50K, basic500K, and basic5000K, respectively. They constituted the training set along with their corresponding stego images. The rest 50% cover-stego pairs in the dataset were for testing. We further guaranteed that the cover images included in an arbitrary training set of the three datasets would not appear in any of the three testing sets.

B. Impact of the framework architecture on the performance

In Tab. I we compare the effect of different Q&T combinations, different hand-crafted convolutional kernels, and the presence of BN layers. The experiment was conducted on basic500K. We can see that under the same conditions, DCT basis patterns (including the 8×8 DCTR kernels [21]) always perform better than PHARM kernels [22]. The experimental results support our choice of DCT basis patterns. 5×5 DCT basis patterns can achieve significant performance improvement compared to 3×3 DCT basis patterns. However, the performance of the more complex 8×8 DCTR kernels is not even as good as the 3×3 DCT basis patterns, which indicates that increasing the size of the convolutional kernels is not always beneficial at the cost of increasing model complexity. Different Q&T combinations also affect the performance of our proposed framework. Combinations with three different quantization steps and the same threshold are of relatively cost-effective. BN layers in the subnets are crucial, especially the first one at the bottom of the subnets. Finally, based on the described results, we adopt twenty-five 5×5 DCT basis patterns, $T = 4, Q = [1, 2, 4]$ and subnet configurations with BN layers following every convolutional layers in our final proposed framework.

C. Comparison to state of the art

In Fig. 4, we compare the performance of our proposed framework and other steganalytic models in the literature. The model proposed by Xu et al. [27] is referred as Xu's model. From Fig. 4 we can see that our proposed framework can obtain significant performance improvement compared with DCTR [21], PHARM [22], and even recently proposed selection-channel-aware JPEG rich model SCA-PHARM [23].

TABLE I
EFFECT OF DIFFERENT Q&T COMBINATIONS, DIFFERENT HAND-CRAFTED CONVOLUTIONAL KERNELS, AND THE PRESENCE OF BN LAYERS. THOSE HYPER-PARAMETERS ADOPTED IN OUR PROPOSED FRAMEWORK ARE MARKED IN BOLD. [a]

Quantization Steps & Threshold	BN Layers		
	With All	Without BN1	Without All
Nine 3×3 DCT basis patterns			
(4,1), (4,1.5), (4,2)	73.1%	70.6%	50.1%
(4,2), (4,2), (4,2)	72.8%	70.1%	50.0%
(4,1), (4,2), (4,4)	<u>73.2%</u>	71.0%	50.1%
(2,1), (4,2), (6,4)	71.2%	68.5%	50.0%
(6,1), (4,2), (2,4)	70.6%	67.8%	50.0%
Twenty-five 5×5 **DCT basis patterns**			
(4,1), (4,1.5), (4,2)	74.3%	72.4%	50.1%
(4,2), (4,2), (4,2)	74.1%	72.4%	50.1%
(4,1)	70.8%	69.4%	50.1%
(4,1), (4,2)	72.5%	70.2%	50.1%
(4,1), (4,2), (4,4)	<u>74.5%</u>	72.5%	50.1%
(2,1), (4,2), (6,4)	73.6%	72.0%	50.1%
(6,1), (4,2), (2,4)	72.6%	71.7%	50.0%
Sixty-four 8×8 DCTR kernels [21]			
(4,1), (4,1.5), (4,2)	72.5%	71.4%	50.0%
(4,2), (4,2), (4,2)	72.7%	71.2%	50.1%
(4,1), (4,2), (4,4)	<u>72.9%</u>	71.2%	50.1%
(2,1), (4,2), (6,4)	71.9%	70.2%	50.0%
(6,1), (4,2), (2,4)	71.5%	70.1%	50.1%
Thirty 5×5 PHARM kernels [22]			
(4,1), (4,1.5), (4,2)	72.0%	70.8%	50.1%
(4,2), (4,2), (4,2)	70.6%	68.8%	50.0%
(4,1), (4,2), (4,4)	<u>72.1%</u>	70.8%	50.1%
(2,1), (4,2), (6,4)	70.3%	68.6%	50.0%
(6,1), (4,2), (2,4)	70.2%	68.7%	50.0%

[a]Logograms are used in expressing Q&T combinations. For example, (4,1) denotes ($T = 4, Q = 1$). The best results in every sub-tables are underlined. Only J-UNIWARD stego images with 0.4bpnzAC were included in the experiments.

For all of the three steganographic algorithms, the performance of Xu's model was poor. The superiority of our proposed framework is more obvious in basic500K. This is due to the fact that with more training samples raised by one magnitude, the large-scale basic500K dataset with 500,000 training samples (covers plus the corresponding stegos) is more favor of deep-learning frameworks like the one proposed by us. If only consider the performance of a single model, our proposed framework with *Type1* subnets behaved better than its companion with *Type2* subnets. Furthermore, the final prediction conducted by the ensemble of five independently trained models shows that model ensemble could improve the detection accuracy by 1% regardless of the type of the underlying subnet configurations. Since the performance of our proposed framework with *Type1* subnets is always better than that with *Type2* subnets, we insisted on using *Type1* subnets in the following experiments. However, please note that *Type2* subnet can potentially be used in more complex deep-learning steganalytic frameworks in the future since it is a memory-efficient model.

In Fig. 5 we show how the testing accuracy changes

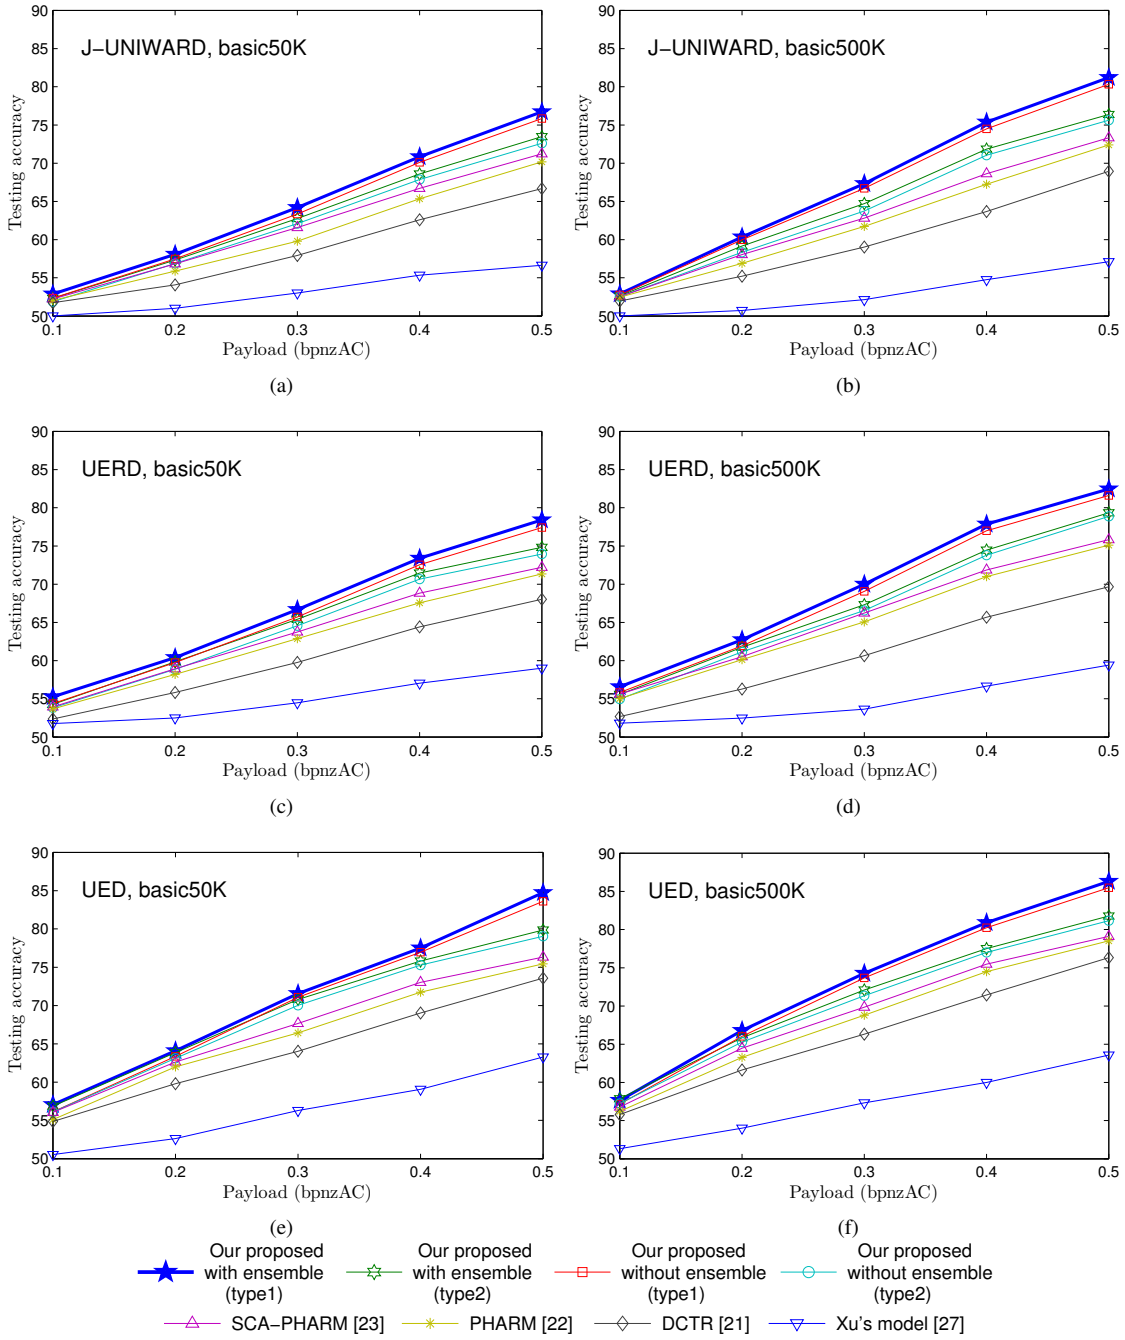

Fig. 4. Comparison of testing accuracy of our proposed frameworks with state-of-the-art steganalytic models described in the literature, two hand-crafted JPEG domain rich models (DCTR and PHARM), a selection-channel-aware variant of PHARM (SCA-PHARM) and a deep-learning steganalytic model proposed by Xu et al. [27]. (a) and (b) are the results for J-UNIWARD; (c) and (d) are for UERD; (e) and (f) are for UED. The experiments for (a), (c) and (e) were conducted on basic50K, while those for (b), (d) and (f) were conducted on basic500K.

with successive training iterations in the experiments which were conducted on basic50K, basic500K and basic5000K, our largest-scale dataset, respectively. The tests were performed on standalone testing dataset every 1×10^4 training iterations and the models were trained for 20×10^4 iterations in total. Only stego images with 0.4bpnzAC were included in the experiments due to the limited computational capacity. Even so for basic5000K there were five million images (covers plus the corresponding stegos) evolved in a training epoch. Our proposed deep-learning framework showed strong learning capacity that further improves along with the growth of training samples. From Fig. 5 we can also see that the curve of testing accuracy for the framework trained on basic5000K not only is of the best performance but also is of the best stability. Please note that 20×10^4 iterations is roughly equivalent to 256 epochs for basic50K, 25.6 epochs for bsic500K, and only 2.56 epochs for basic5000K. Therefore the full potential of our proposed framework with large-scale training datasets may not

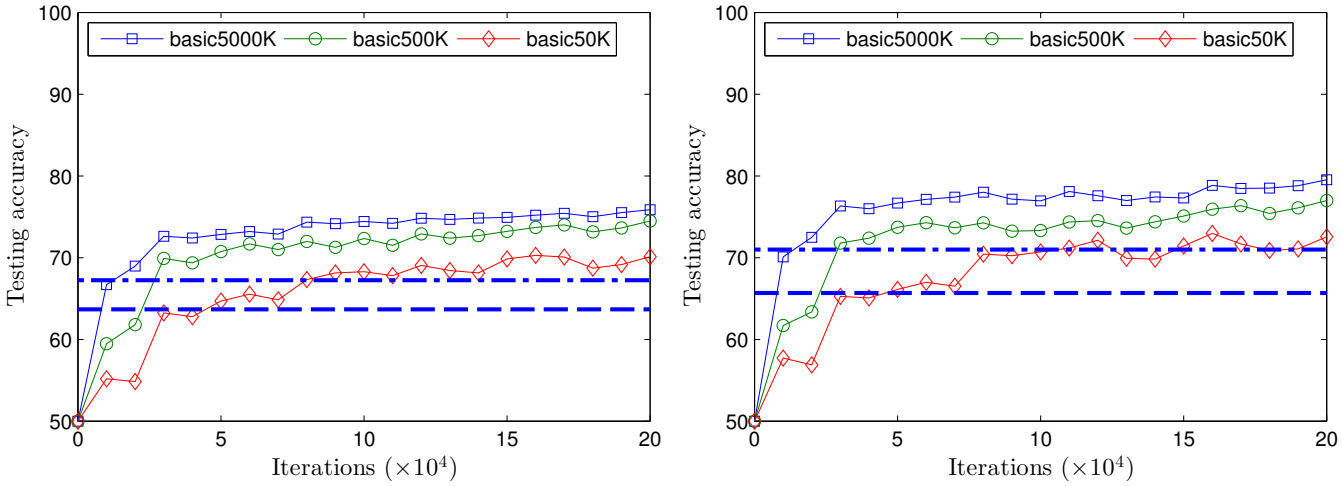

Fig. 5. Testing accuracies versus training iterations for our proposed framework. The experiments were conducted on basic50K, basic500K and basic5000K, respectively. Only stego images with 0.4bpnzAC were included in the experiments. The dash-dotted reference line denotes the best testing accuracy of PHARM in basic500K, while the dashed reference line denotes the best testing accuracy of DCTR in basic500K. (a) J-UNIWARD steganography. (b) UERD steganography.

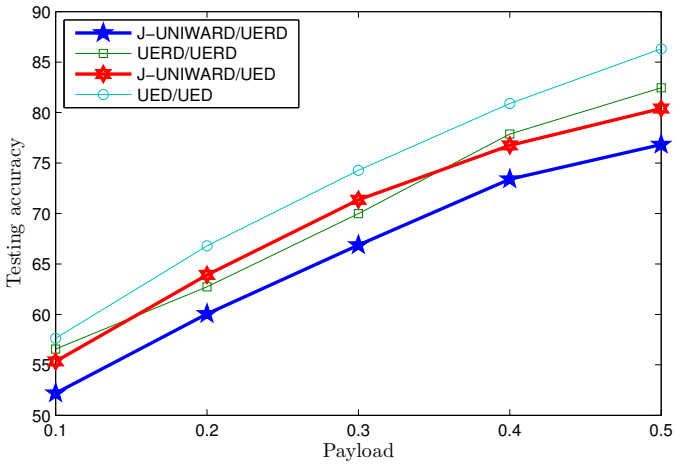

Fig. 6. Comparison of attacking-target transfer ability of our proposed framework. The experiments were conducted on basic500K dataset. Only stego images with 0.4bpnzAC were included in the experiments. The notations in the legend take the form of the target in training and the target in testing delimited by a slash (/). For example, "J-UNIWARD/UERD" means that J-UNIWARD stego images were used in training while UERD stego images were used in testing.

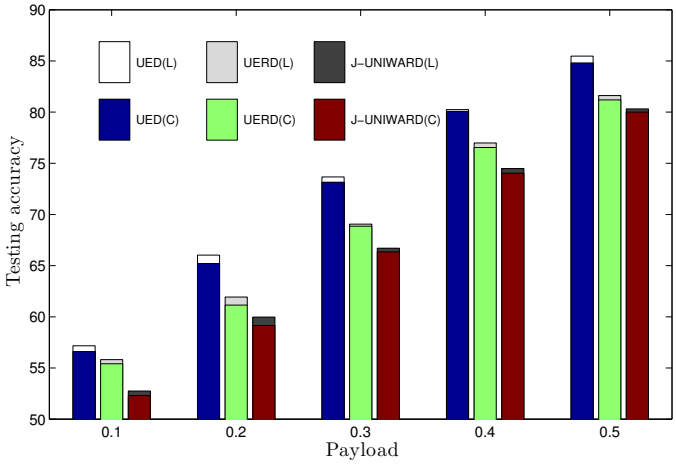

Fig. 7. The impact of altered blocking artifact on the performance of our proposed framework. Only stego images with 0.4bpnzAC were included in the experiments. All of the models were trained on basic500K training set and then tested on the corresponding testing set with central-cropped images. The legend "C" in parentheses denotes those tested on central-cropped images, while "L" in parentheses denotes those tested on the original basic500K testing set. For example, "J-UNIWARD (C)" means that the corresponding framework was trained and tested with J-UNIWARD stego images. It was trained on basic500K training set and then tested on the corresponding testing set with central-cropped images.

have been fully exploited. [4]

Throughout the experiments, our proposed framework ran steadily. During the training procedure, it could accomplish 1,000 iterations every 20 minutes. That is to say, 20×10^4 training iterations could be finished in about 67 hours. With K80 GPU cards, We can expect to finish one epoch of training in 0.26 hour, 2.6 hours, and 26 hours for basic50K, basic500K, and basic5000K, respectively.

[4]The implementation of ensemble classifier [20] used by rich models cannot be scaled to large-scale datasets. Therefore we cannot provide the testing accuracy of DCTR and PHARM in basic5000K for comparison in Fig. 5.

D. Performance with mismatched targets, altered blocking artifact and doubled-sized input images

In Fig. 6, we observe the attacking-target transfer ability of our proposed framework. The framework was trained with J-UNIWARD cover/stego pairs and then tested with UERD/UED cover/stego pairs. The detection accuracy is roughly $3\% - 4\%$ worse compared with that trained and tested with the same type of stego images. However, the degradation of detection performance is acceptable especially for the detection of UED stego images given that UED works in a very different way compared with J-UNIWARD.

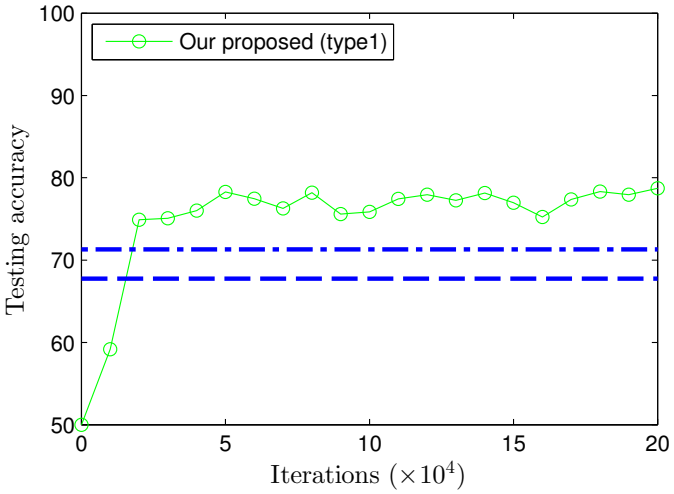

Fig. 8. Testing accuracies versus training iterations for our modified framework which takes 512×512 images as input. Only J-UNIWARD stego images with 0.4bpnzAC are included in the experiment. As in Fig. 5, The dash-dotted reference line denotes the best testing accuracy of PHARM, while the dashed reference line denotes the best testing accuracy of DCTR in the same testing dataset.

8×8 block processing during JPEG compression introduces blocking artifacts, which can be used as intrinsic statistical characteristic of JPEG cover images. Secret bits embedded in the DCT domain tend to impair blocking artifacts, therefore leave traces which can be utilized by steganalyzers. An interesting problem is to access the performance of our proposed framework depending on the intrinsic statistical characteristic of blocking artifacts. In Fig. 7, we observe the impact of altered blocking artifacts on the performance of our proposed framework. The default testing set in basic500K contains left-top cropped images in which the original DCT grid alignment is preserved. In this experiment for every testing images in basic500K, we re-compressed their corresponding original images in ImageNet with quality factor 75 and then converted them to grayscale images again. We cropped their central 256×256 regions to constitute a new testing set. The motivation is that central cropping cannot preserve the original DCT grid alignment in most cases. Double compression preserving the original DCT grid alignment implies that the existing blocking artifacts in the original ImageNet images are depressed by the blocking artifacts introduced in double compression. Double compression not preserving the original DCT grid alignment, by contrast, implies that the blocking artifacts from two different sources coexist. As a result the blocking artifacts in the images of the new testing set are different from those in the training set. However, from Fig. 7 we can see the performance of our proposed framework on the new testing set was almost identical with the original testing set in basic500K. The experimental results reveal that the impact of altered blocking artifact on the performance of our proposed framework is small. Our proposed framework has captured more complex intrinsic statistical characteristic besides blocking artifact.

Most of the conducted experiments used images of size 256×256 pixels. This limitation stems mainly from the following two factors: Firstly, the scale of deep-learning models is still constrained by GPU memory. At the current stage, target images with larger size, e.g. 512×512 pixels result in deep-learning models hard to train with K80 GPU cards we have in hands. Secondly, large-sized ImageNet images are in the minority. Out of fourteen million ImageNet images, only roughly 0.7 million of them are larger than 512×512 pixels. In the last experiment, we tested our proposed framework on this limited dataset. 500 thousand JPEG images with size larger than 512×512 were randomly picked out from ImageNet and were converted to 512×512 with the same processing procedure as mentioned in Sect. III-A. Due to GPU memory constraints, we simplified the model by using doubled stride in the convolutional layer of each subnet (i.e. 4 instead of 2). All other experiment setups were remained the same except that the batch size in the training procedure is reduced to 32. Only J-UNIWARD stego images with 0.4bpnzAC were included in the experiment. Fig. 8 shows the testing accuracy in successive training iterations. Compared to that for 256×256 images (Fig. 8), the simplified framework can get better detection performance. The training procedure converged quickly and delivered better performance than the DCTR and PHARM models. Due to the limited computational capacity, subnets with wider and deeper structures were not evaluated in this last experiment. Its potential for target images with larger size may have not been fully demonstrated.

IV. Concluding remarks

Application of deep-learning frameworks in image steganalysis has drawn attention of many researchers. In this paper we proposed a hybrid deep-learning framework for large-scale JPEG steganalysis. The major contributions of this work are as follows:

- We have proposed a hybrid deep-learning framework for JPEG steganalysis, which contains a hand-crafted stage incorporated from rich models and a learnable stage composed of a compound deep CNN network.
- We have proved that the convolution phase and the Q&T phase in rich models are not learnable in deep-learning framework, as the theoretical basis of the hand-crafted stage of our proposed framework.
- We have conducted extensive experiments on a large-scale dataset extracted from ImageNet. The largest dataset used in our experiments contains ten million covers and stegos. This is the first report of deep-learning steganalytic frameworks evaluated on large-scale datasets.

Our future work will focus on two aspects: (1) incorporation of adversarial machine learning into our proposed framework to make it jointly optimized with its opponent; (2) further exploration of the application of our proposed framework in the field of multimedia forensics.

Acknowledgment

The authors would like to thank DDE Laboratory in SUNY Binghamton for sharing the source code of their steganalysis models online (http://dde.binghamton.edu/download/). We also appreciate Prof. Jiangqun Ni in Sun Yat-sen University, China

for permission to use their implementation of UED and UERD in our experiments. Specifically, we are grateful to Dr. Paweł Korus in Shenzhen University for valuable advice.

References

[1] I. Cox, M. Miller, J. Bloom, J. Fridrich, and T. Kalker, *Digital Watermarking and Steganography*, 2nd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2008.

[2] M. S. Subhedar and V. H. Mankar, "Current status and key issues in image steganography: A survey," *Computer Science Review*, vol. 13, pp. 95–113, 2014.

[3] J. Fridrich and T. Filler, "Practical methods for minimizing embedding impact in steganography," in *Proc. SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX*, vol. 6505, 2007, pp. 650 502–1–650 502–15.

[4] T. Pevný, T. Filler, and P. Bas, "Using high-dimensional image models to perform highly undetectable steganography," in *Proc. 12th Information Hiding Workshop (IH'2010)*, 2010, pp. 161–177.

[5] V. Holub and J. Fridrich, "Designing steganographic distortion using directional filters," in *Proc. IEEE 2012 International Workshop on Information Forensic and Security (WIFS'2012)*, 2012, pp. 234–239.

[6] B. Li, S. Tan, M. Wang, and J. Huang, "Investigation on cost assignment in spatial image steganography," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 8, pp. 1264–1277, 2014.

[7] B. Li, M. Wang, J. Huang, and X. Li, "A new cost function for spatial image steganography," in *Proc. IEEE 2014 International Conference on Image Processing, (ICIP'2014)*, 2014, pp. 4206–4210.

[8] J. Fridrich and J. Kodovský, "Multivariate Gaussian model for designing additive distortion for steganography," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'2013)*, 2013, pp. 2949–2953.

[9] V. Sedighi, J. Fridrich, and R. Cogranne, "Content-adaptive pentary steganography using the multivariate generalized Gaussian cover model," in *Proc. IS&T/SPIE Electronic Imaging 2015 (Media Watermarking, Security, and Forensics)*, 2015, pp. 94 090H–1–94 090H–13.

[10] V. Sedighi, R. Cogranne, and J. Fridrich, "Content-adaptive steganography by minimizing statistical detectability," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 2, pp. 221–234, 2016.

[11] L. Guo, J. Ni, and Y. Q. Shi, "Uniform embedding for efficient JPEG steganography," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 5, pp. 814–825, 2014.

[12] L. Guo, J. Ni, W. Su, C. Tang, and Y. Q. Shi, "Using statistical image model for JPEG steganography: Uniform embedding revisited," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 12, pp. 2669–2680, 2015.

[13] V. Holub, J. Fridrich, and T. Denemark, "Universal distortion function for steganography in an arbitrary domain," *EURASIP Journal on Information Security*, vol. 2014, no. 1, pp. 1–13, 2014.

[14] T. Denemark and J. Fridrich, "Improving steganographic security by synchronizing the selection channel," in *Proc. 3rd ACM Information Hiding and Multimedia Security Workshop (IH&MMSec'2015)*, 2015, pp. 5–14.

[15] B. Li, M. Wang, X. Li, S. Tan, and J. Huang, "A strategy of clustering modification directions in spatial image steganography," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 9, pp. 1905–1917, 2015.

[16] J. Fridrich and J. Kodovský, "Rich models for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 3, pp. 868–882, 2012.

[17] W. Tang, H. Li, W. Luo, and J. Huang, "Adaptive steganalysis against WOW embedding algorithm," in *Proc. 2nd ACM Information Hiding and Multimedia Security Workshop (IH&MMSec'2014)*, 2014, pp. 91–96.

[18] T. Denemark, V. Sedighi, V. Holub, R. Cogranne, and J. Fridrich, "Selection-channel-aware rich model for steganalysis of digital images," in *Proc. 6th IEEE International Workshop on Information Forensic and Security (WIFS'2014)*, 2014, pp. 48–53.

[19] W. Tang, H. Li, W. Luo, and J. Huang, "Adaptive steganalysis based on embedding probabilities of pixels," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 4, pp. 734–745, 2016.

[20] J. Kodovský and J. Fridrich, "Ensemble classifiers for steganalysis of digital media," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 2, pp. 432–444, 2012.

[21] V. Holub and J. Fridrich, "Low-complexity features for JPEG steganalysis using undecimated DCT," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 2, pp. 219–228, 2015.

[22] ——, "Phase-aware projection model for steganalysis of JPEG images," in *Proc. IS&T/SPIE Electronic Imaging 2015 (Media Watermarking, Security, and Forensics)*, 2015, pp. 94 090T–1–94 090T–11.

[23] T. Denemark, M. Boroumand, and J. Fridrich, "Steganalysis features for content-adaptive JPEG steganography," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 8, pp. 1736–1746, 2016.

[24] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Networks*, vol. 61, pp. 85–117, 2015.

[25] S. Tan and B. Li, "Stacked convolutional auto-encoders for steganalysis of digital images," in *Proc. Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA'2014)*, 2014.

[26] Y. Qian, J. Dong, W. Wang, and T. Tan, "Deep learning for steganalysis via convolutional neural networks," in *Proc. IS&T/SPIE Electronic Imaging 2015 (Media Watermarking, Security, and Forensics)*, 2015, pp. 94 090J–1–94 090J–10.

[27] G. Xu, H. Z. Wu, and Y. Q. Shi, "Structural design of convolutional neural networks for steganalysis," *IEEE Signal Processing Letters*, vol. 23, no. 5, pp. 708–712, 2016.

[28] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv:1502.03167*, 2015. [Online]. Available: http://arxiv.org/abs/1502.03167

[29] P. B. T. Filler and T. Pevný, "Break our steganographic system—the ins and outs of organizing BOSS," in *Proc. 13th Information Hiding Workshop (IH'2011)*, 2011, pp. 59–70.

[30] V. Sedighi, J. Fridrich, and R. Cogranne, "Toss that BOSSbase, Alice!" in *IS&T/SPIE Electronic Imaging (Media Watermarking, Security, and Forensics)*, San Francisco, CA, USA, 14-18 February 2016.

[31] "ImageNet," http://image-net.org/, [Online].

[32] "CS231n: Convolutional Neural Networks for Visual Recognition," http://cs231n.github.io/, [Online].

[33] B. Bayar and M. C. Stamm, "A deep learning approach to universal image manipulation detection using a new convolutional layer," in *Proc. 4th ACM Information Hiding and Multimedia Security Workshop (IH&MMSec'2016)*, 2016, pp. 5–10.

[34] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1–9.

[35] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv:1408.5093*, 2014. [Online]. Available: http://arxiv.org/abs/1408.5093